

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE
APPLICATION FOR U.S. LETTERS PATENT

Title:

AN OPTICAL INTERCONNECT IN HIGH-SPEED MEMORY SYSTEMS

Inventors:

Jacob Baker
Brent Keeth

Thomas J. D'Amico
DICKSTEIN SHAPIRO MORIN &
OSHINSKY LLP
2101 L Street NW
Washington, DC 20037-1526
(202) 828-2232

FIELD OF THE INVENTION

[0001] The present invention relates to communicating at high speed data signals to and from memory storage devices such as DRAM memory devices.

BACKGROUND OF THE INVENTION

[0002] As computer processor and DRAM (Dynamic Random Access Memory) memory speeds increase, their bus speeds increase also. This increased speed also increases signal noise at connection points where a memory controller and DRAM memory devices connect to a bus. In addition, the connections of the bus also have associated electrical properties such as capacitance and inductance which, while causing minimal problems at low data speeds, causes increasingly significant problems at high speed. Consequently, at high speed, conventional bus arrangements can introduce signal distortion, noise, delays and other unwanted spurious signal phenomenon.

[0003] Current memory devices commonly operate at hundreds of megahertz, but it is anticipated that computer bus speeds, which tend to run slightly slower than microprocessor speeds, will soon extend beyond 1GHz. At such high frequencies, the minutest amount of signal aberration caused by the electrical properties of the electrical bus may cause severe and unexpected consequences. Additionally, the distance between components on a bus must be kept short, to minimize signal distortions and help insure that data and control signals reach their destination very quickly.

[0004] Accordingly, a memory bus structure which reduces or eliminates signal distortion, noise, and other problems and permits reliable high speed (e.g. greater than 1 GHz) operation is desired.

BRIEF SUMMARY OF THE INVENTION

[0005] In one aspect the invention provides a memory apparatus and method of its operation which utilizes an optical path connected between a memory controller or processor and at least one memory device for passing data between the controller or processor and memory device at high throughput speed.

BRIEF DESCRIPTION OF THE DRAWINGS

[0006] The foregoing and other features and advantages of the invention will become more apparent from the detailed description of the exemplary embodiments of the invention given below with reference to the accompanying drawings in which:

Fig. 1 shows a generic overview of the present invention;

Fig. 2 shows one exemplary embodiment of the invention;

Fig. 3 shows a transistor-level view of the transmitter and receiver used in an exemplary embodiment of the invention;

Fig. 4 shows a second exemplary embodiment of the invention;

Fig. 5 shows a third exemplary embodiment of the invention;

Fig. 6 shows a fourth exemplary embodiment of the invention;

Fig. 7 shows a fifth embodiment of the invention.

DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENTS

[0007] The present invention uses one or more optical links between a processor and/or a memory controller and a DRAM memory device. The optical link includes, but is not limited, to optical fiber and optical waveguide links as described below in connection with various exemplary embodiments of the invention. Fig. 1 shows a high level block diagram of the present invention. A processor 100 is connected to a memory controller 104 which in turn is connected to a memory module 113 containing one or more memory devices 112 using one or more optical links 108. The memory controller 104 and modules 113 have optical couplers which enable them to connect to the optical links 108 to maintain optical continuity. The modules 113 have optical plug-in connectors to the optical links 108, but also have standard (non-optical) Dual Inline Memory Module (DIMM) connectors 109 for supplying power and other low-frequency signals.

[0008] In the context of the invention, the processor 100, controller 104, and memory devices 112 can be located either on the same die or located on separate dies. In some cases, processor 100 can also serve as the memory controller 104 in which case a separate memory controller 104 can be omitted.

[0009] Fig. 2 shows a first exemplary embodiment of the invention in which a single common optical link 108a transmits a plurality of data streams between a memory controller 104 and memory modules 113 using paired optical transmitters and receivers on opposite sides of link 108a pre-set to communicate at a respective wavelength. Fig. 2 shows the use of separate data (DQ), command (CMD), address (ADD), and clock (CLK) paths between controller 104 and each memory module 113 as is typical in a computer bus structure. It is also possible to send control and address data over the same data paths as is also well known in the art. For brevity, only the data (DQ) optical path will be discussed in detail, it being understood that the optical paths for other data and clock information sent

by the controller will be handled the same except for the direction of data/clock pulse flow. It should also be understood that while the data (DQ) paths are bidirectional, the command/address and clock paths are unidirectional in that the dataflow is from controller 104 to the modules 113 and associated memory devices 112.

[0010] As shown in Fig. 2, each data DQ path of the memory controller 104 is coupled to a respective optical transmitting/receiving device $T_0/R_0 \dots T_{15}/R_{15}$, each collectively identified by the label 201. Each transmitting/receiving device converts an electrical signal received from a DQ path of memory controller 114 and converts the electrical signal to an optical signal for transmission on optical link 108a to a memory module 113 over optical link 108a. Each transmitter/receiver 201 is also capable of receiving an optical signal from a module 113 and converting it to an electrical signal and sending it to controller 104 on a respective data (DQ) path.

[0011] In addition to the transmitter/receivers 201 provided on the controller side, respective transmitters 203 are also provided for converting each of the electrical signals on the command, address and clock signal paths to optical signals over link 108a and transmitting these optical signals to modules 113. The transmitter/receivers 201 and transmitters 203 may form part of an electrical/optical converter 205.

[0012] The Fig. 2 embodiment uses a single optical link 108a constructed as an optical fiber or optical waveguide between controller 104 and the memory modules 113. In this way, many datapins of controller 104 communicate over a single optical link 108a. In order to keep the optical signals from the different data (DQ), command (CMD), address (ADDRESS), and clock (CLK) paths from interfering with each other, wave division multiplexing is employed so that the optical signals from each of the transmitter/receiver devices 201 and transmitter devices 203 have a respective optical

carrier wavelength (frequency) which is modulated by data sent on the various signal paths from controller 104 to converter 205. Likewise, the optical receiver portion of each transmitter/receiver 201 operates at a respective optical wavelength.

[0013] As further shown in Fig. 2, the various optical signals from transmitter/receivers 201 and transmitters 203 are optically combined in a multiplexing portion of a wavelength division multiplexer/demultiplexer 207 for transmission over the common optical link 108a to memory modules 113.

[0014] Each module 113 also contains a wave division multiplexer/demultiplexer 209 which receives the optically multiplexed signals on optical link 108a and wavelength demultiplexes them in a demultiplexer portion and passes the demultiplexed signals to respective transmitter/receivers 211, which electrically connect to the data (DQ) paths of the memory devices 112. In addition, the demultiplexed optical signals for the command (CMD), address (ADD) (or combined command/address) and clock (CLK) signal paths are passed on to receivers 213 which convert optical signals to electrical signals which are electrically coupled to the electrical command (CMD), address (ADD) and clock (CLK) signal paths of the memory devices 112.

[0015] Data read from memory devices 112 is transmitted on the data (DQ) paths of the memory devices 112 to respective transmitter/receivers 211 where the electrical data is converted to an optical signal at a respective wavelength and sent to multiplexer/demultiplexer 209 where the data on the respective DQ optical paths is combined in the wave division multiplexer of multiplexer/demultiplexer 209. This data is then sent over optical link 108a to multiplexer/demultiplexer 207 where it is demultiplexed and passed to respective transmitter/receivers 201 where the DQ optical data is connected to electrical DQ data which is sent to respective DQ data paths of

controller 104. Fig. 2 illustrates the optical coupling of two memory modules 113 to memory controller 104 through the electro-optical converter 205 provided at the memory controller 104 side of optical link 108 and an electro-optical converter 219 provided on the memory modules 113; however, it should be understood that any number of memory modules 113, containing any number of memory devices 112, may be optically coupled to controller 104 over optical link 108a.

[0016] Fig. 3 shows a simplified optical transmitter 116 and optical receiver 120 which may be used in the electro/optical transmitter/receivers 201, 211 and in the electro/optical transmitters 203 and receivers 213. A LED (Light Emitting Diode) or ILD (Injection Laser Diode) light emitter 124 in transmitter 116 provides a light output signal to an optical path 241 at a predefined wavelength, in response to an applied electrical signal at the gate of a transistor 126. At the receiver 120 side, a photodiode 128 couples light pulses received from an optical path 241 to the gate of an n-channel transistor 134. A p-channel biasing transistor 138 sources current to the n-channel transistor 134. A resistor 135 is positioned between the gate of transistor 134, as well as the drain of transistor 138. The transistors 134 and 138 and resistor 135 form an inverting amplifier 137. The output 139 of the inverting amplifier 137 is an electrical signal.

[0017] Although Fig. 3 illustrates the light transmitter 116 and receiver 120 as discrete components, these devices are actually integrated devices which may be integrated together with multiplexer/demultiplexer 207 on a converter 205 chip or integrated on the same chip as the memory controller 104. At the module 113, the transmitter 116 and receiver 120 are preferably integrated on the same chip which contains the multiplexer/demultiplexer 209. It is also possible to integrate the transmitter 116 and

receiver 120 on the module side within the actual memory devices 112 in which case each memory device 112 would contain its own converter circuit 219 shown in Fig. 3.

[0018] Although a silicon substrate may be used for integrating the LED or ILD light emitter 124 and/or photodiode 128, the more preferred substrate material for such devices, particularly for LED or ILD 124 is gallium arsenide, as known in the art. Finally, it should be understood that while Fig. 3 illustrates a unidirectional data path, in actuality the data (DQ) paths in a memory system are bi-directional and that an optical transmitter 116 and receiver 120 are therefore understood to be employed at each path end of a bidirectional optical link 108a, as shown by transmitter/receivers 201 and 211.

[0019] As noted, the Fig. 2 arrangement relies on wavelength division multiplexing of the different signal paths which exist between memory controllers 104 and the individual memory devices 112. Thus, each transmitter/receiver 201, transmitter 203 and receiver 235 as well as multiplexer/demultiplexers 207, 209 must operate at specified optical wavelengths. These wavelengths can be controlled using known filter circuits. However, it is often difficult to ensure that a manufacturer's device operates precisely at a predetermined wavelength. To this end, it is also known to adjust operating conditions of an electro/optical device to ensure that it operates at a predetermined wavelength.

[0020] Fig. 4 shows a modification of a portion of the system of Fig. 2, where transmitting devices 201 and receiving devices 203 are shown as being wavelength-adjustable. For clarity, only the DQ0 pin is shown, while DQ1-DQ15 are implied, similar to the representation in Fig. 2. During fabrication, the thicknesses and purities of the materials deposited as well as other factors make it difficult to fabricate a transmitter 203 and the transmitter portion of receiver/transmitters 201 and 211 to transmit at a precise predefined wavelength. Accordingly, the light emitters are wavelength adjustable.

Wavelength detectors 233 are used to sense the nominal wavelength of an optically transmitted signal from each of the transmitters of devices 201 and 203 and data representing the sensed wavelength is fed back to controller 104 which determines if a transmitter is transmitting at its assigned wavelength and, if not, a wavelength adjuster 231 is operated by controller 104 which sends data to an addressed wavelength adjuster 231 for adjusting the wavelength over the command (CMD) signal path. Separate control signal paths can also be used for this purpose. The wavelength of optical signals sent by the data transmitters 211 in the modules 113 can also be sensed by the wavelength detector 233 and adjustment data can be sent to addressed wavelength adjuster 235 on the module 113 which adjusts the wavelength of the transmitter portion of transmitter/receiver 211. The adjustments can be accomplished during initialization of the memory system for operation.

[0021] Fig. 5 shows another embodiment of the invention, which utilizes an optical link 108b for each data path on an optical bus 111. In this embodiment there is a one-to-one replacement of an electrical bus line which normally interconnects memory controller 104 with a memory module 113 with an optical link 108b. For simplicity, Fig. 5 only shows four such optical links (two DQ, one CMD of a CLK path). The individual optical links 108b connect with transmitter/receivers 211 or receivers 213 on the memory modules which convert the optical signals to electrical signals for use by memory devices 112 and electrical signals to optical signals for data read from the memory devices 112.

[0022] As seen, there are several different techniques of optical data transmission which can be used on the optical link 108 in the present invention. These techniques can include but are not limited to Time Division Multiplexing (TDM). Using TDM, data from multiple pins can be used to occupy a single optical channel. Also, TDM can be used in

conjunction with other optical data transmission schemes to reduce the number of optical channels (either fiber or wavelength) needed within an optical system. Two more examples of such techniques are Wavelength Division Multiplexing (WDM) and Frequency Division Multiplexing (FDM). Additionally, data compression techniques can be used. Such techniques have in common that they reduce the volume of data transmitted, the number of optical channels needed, or both.

[0023] An embodiment of the present invention using WDM is shown in Fig. 2. WDM enables the simultaneous transmission of multiple data channels on the same physical optical link, by utilizing several different wavelengths on that optical link at the same time. An optical multiplexer (mux) portion of the multiplexer/demultiplexer 207, 209 combines different wavelength bands from individual optical sources into a multiple wavelength light beam for simultaneous transmission through a common optical link. At the receiving end of the optical link, an optical demultiplexer (demux) portion of a multiplexer/demultiplexer 209 demultiplexes or spatially disburses collimated multiple wavelength light from the optical link into separate wavelength bands, each of which can be directed to an individual optical receiver. Although Fig. 2 shows combination of multiplexer/demultiplexer devices 207, 209 it should be apparent that separate multiplexers and demultiplexers can be used as well to perform the required multiplexing and demultiplexing functions. Another optical transmission technique, as shown in Fig. 5, uses a separate optical link for each data path.

[0024] It should also be noted that although all data paths (e.g., write/read data (DQ), command (CMD), address (ADD), clock (CLK) between the memory controller 104 and modules 113 are shown as utilizing optical transmission, it is also possible to use optical transmission only on the high speed data paths, e.g. the write/read data (CD) and

clock (CLK) paths and utilize conventional electrical bus lines for slower speed data paths, e.g. command (CMB), address (ADD).

[0025] The present invention can use any modulation format in the optical link to optimize either Signal to Noise Ration (SNR) or bandwidth utilization. This could include conventional digital modulation techniques such as FM or Non Return To Zero (NRTZ).

[0026] The processor 100, controller 104, and memory devices 112 are typically located on separate dies with the memory devices being mounted on modules 113 which connect with the optical link 108a or 108b. However, it is also possible to integrate the processor and memory devices on the same die, with the processor incorporating the functions of the memory controller or with the memory controller also being integrated on the processor die. In the case where they are located on the same die, an integrated optical waveguide can be used to link them. Fig. 6, for example, shows an exemplary confined square pipe waveguide 212. Positioned on die 200, the waveguide 202 connects a processor with an integrated memory controller with DRAM 112. The waveguide 200 has a first metal layer 208 on top, a second metal layer 210 on the bottom, end plates 212 connecting the top and bottom layers, and an optically transmissive insulator 214 in middle through which light pulses carrying data are transmitted. The two metal layers (208, 210) act as waveguides confining the light pulses. The insulator 214 could be made of SiO₂ which is commonly used in chip formation. Furthermore, in those configurations where the processor 204 and memory devices 206 are not on the same wafer or die and the module 113 and controller 104 are omitted, the waveguide 202 could also be implemented in freespace (air or vacuum).

[0027] Fig. 7 shows an optical link 108c in the form of a flexible optical fiber. Using such a fiber, a processor 100 and memory devices 112 can be integrated on separate dies

residing in separate planes and packaged separately or together, with the processor 100 and memory devices 112 being interconnected by the flexible optical fiber 108c. This allows easier fabrication of the bus lines as well as non-planar stacking of processor 100 and DRAM devices 112 in separate or common packaging.

[0028] All of the above embodiments have in common that they achieve electrical isolation between the memory device 112 and the controller 104. They also make the optical link 108a, 108b, and 108c interconnections immune to noise, including at high frequency. Because the link is operated at high frequency, the clock signal for latching in data is sent with the data. Because fiber optic links do not affect pulse shape as do conventional electrical links, the memory devices 112 can be placed a greater distance from the controller 104 than is conventional. An additional advantage of the invention is that fiber optic links have lower power dissipation than conventional electrical links. This is because fiber optic links do not require I/O buffers, which consume power and also slow the propagation rate at which data is transferred.

[0029] While the invention has been described and illustrated with reference to specific exemplary embodiments, it should be understood that many modifications and substitutions can be made without departing from the spirit and scope of the invention. Accordingly, the invention is not to be considered as limited by the foregoing description but is only limited by the scope of the appended claims.